

L Number	Hits	Search Text	DB	Time stamp
1	34	leader\$5 with elect\$6 with (multicast\$5 group\$6)	USPAT	2003/09/19 15:36
2	10	(y near huang).in.	USPAT	2003/09/19 15:37
3	259	singh near s) (kurose near j	USPAT	2003/09/19 15:38
4	34	(singh near s) (kurose near j) and (7\$/\$.ccls.	USPAT	2003/09/19 15:42
5	1	5738732.pn.	USPAT	2003/09/19 15:42
6	0	538732.pn.	USPAT	2003/09/19 15:42
7	1	5938732.pn.	USPAT	2003/09/19 15:43
8	1	5938732.pn. and (leader\$5 same (elect\$5 select\$5))	USPAT	2003/09/19 15:44
9	1	5938732.pn. and (leader\$5)	USPAT	2003/09/19 15:44
10	0	5938732.pn. and (leader\$5 with priori\$6)	USPAT	2003/09/19 15:44
11	0	5938732.pn. and (leader\$5 same prior\$6)	USPAT	2003/09/19 15:45
12	1	5938732.pn. and (leader\$5 prior\$6)	USPAT	2003/09/19 15:45
13	1	5938732.pn. and (leader\$5 priorit\$6)	USPAT	2003/09/19 15:45
14	1	5699501.pn. and (membership near list\$5)	USPAT	2003/09/19 17:09

[Advanced Search](#) [Preferences](#) [Language Tools](#) [Search Tips](#)

multicast elect leader

Google Search

[Web](#) - [Images](#) - [Groups](#) - [Directory](#) - [News](#)Searched the web for **multicast elect leader**.Results **1 - 10** of about **892**. Search took **0.14** seconds**[PDF] An Analysis of Leader Election for Multicast Groups**File Format: PDF/Adobe Acrobat - [View as HTML](#)... lead router to track host interest in **multicast** groups [6 ... all processes and its results in a single **leader**. ... an algorithm could be used to **elect** multiple leaders ...www.cs.caltech.edu/~schooler/papers/le_analysis.pdf - [Similar pages](#)**Leader Selection**... of the hierarchy is divided into groups formed by the scoping of messages sent out on the **multicast** layer. Within a group, peer agents must **elect a leader**. ...www.cs.caltech.edu/~schooler/marychan/paper/node11.html - 6k - [Cached](#) - [Similar pages](#)[[More results from www.cs.caltech.edu](#)]**Citation details: Ad Hoc On-Demand Distance Vector - Royer, Das, ...**... The **multicast** operation of the Ad hoc On Demand Distance Vector (AODV) routing protocol [6, 7] performs **leader** election to **elect** a new **multicast** group **leader** ...citeseer.nj.nec.com/context/1243593/0 - 6k - [Cached](#) - [Similar pages](#)**Citations: Lightweight causal and atomic group multicast - Birman ...**... decision with a single message delay [1]. If the **leader** fails, the ... eg to remove a faulty primary and **elect** a new ... Lightweight causal and atomic group **multicast**. ...citeseer.nj.nec.com/context/3056/0 - 39k - [Cached](#) - [Similar pages](#)[[More results from citeseer.nj.nec.com](#)]**[PS] 1**File Format: Adobe PostScript - [View as Text](#)... first one to make a willing broadcast is the **leader**; ... primary Primary crashes Eliminate Primary; **Elect** new primary Distributing ... Primary crashes During **multicast**. ...homepages.cs.ncl.ac.uk/paul.ezhilchelvan/home/formal/papers/BloreHandouts.ps - [Similar pages](#)**FoRK Archive: Re: Review requested: Self-Organizing Multicast ID**... Self-Organizing **Multicast**, > MTP/SO ... to build similar large groups that are not **leader**-based, but ... the protocols required to detect their failure, **elect** a new ...www.xent.com/FoRK-archive/feb98/0296.html - 12k - [Cached](#) - [Similar pages](#)**[PPT] CSE 536**File Format: Microsoft Powerpoint 97 - [View as HTML](#)... Ordered according to timestamp. **Multicast** acknowledgement. ... As controller or coordinator.We need to **elect** a **leader** on startup and when current **leader** fails. ...www.eas.asu.edu/~cse536s2/lectures/lect10.ppt - [Similar pages](#)**[PDF] Self Organized Subnets (SOS)**File Format: PDF/Adobe Acrobat - [View as HTML](#)... update (ID, subnet prefix) to Source Node yes no Packet to all-routers **multicast** address no ... Algorithm • Routers **elect** a **leader** • **Leader** collects state ...www.scs.carleton.ca/~adhocnow/NETWORKS/infonet-030325.pdf - [Similar pages](#)**[PDF] An Analysis of Fault Isolation in Multi-Source Multicast Session**File Format: PDF/Adobe Acrobat - [View as HTML](#)... The way to **elect** a group **leader** is, first new join member send a "i-am-leader" packet to subnet using **multicast** scoping not to forward beyond its group, if ...cosmos.kaist.ac.kr/~hkpark/data/hkparkapan0308v2.pdf - [Similar pages](#)

[PDF](#) [A Probabilistically Correct Leader Election Protocol for Large ...](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

... The initiating **multicast** is **multicast** to the entire group, but ... failures, each connected component of this graph will **elect** exactly one **leader**, with a ...

www.cs.cornell.edu/gupta/Papers/disc2000.final.pdf - [Similar pages](#)

Google

Result Page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [Next](#)

[Search within results](#)

Dissatisfied with your search results? [Help us improve.](#)

Get the [Google Toolbar](#):



[Google Home](#) - [Advertise with Us](#) - [Business Solutions](#) - [Services & Tools](#) - [Jobs, Press, & Help](#)

©2003 Google

Welcome to IEEE Xplore®

- ☐ Home
- ☐ What Can I Access?
- ☐ Log-out

Tables of Contents

- ☐ Journals & Magazines
- ☐ Conference Proceedings
- ☐ Standards

Search

- ☐ By Author
- ☐ Basic
- ☐ Advanced

Member Services

- ☐ Join IEEE
- ☐ Establish IEEE Web Account
- ☐ Access the IEEE Member Digital Library

Print Format

Your search matched **5** of **971567** documents.A maximum of **5** results are displayed, **15** to a page, sorted by **Relevance** in **descending** order.

You may refine your search by editing the current search expression or entering a new one in the text box.

Then click **Search Again**.

Results:

Journal or Magazine = **JNL** Conference = **CNF** Standard = **STD****1 Group leader election under link-state routing***Yih Huang; McKinley, P.K.;*

Network Protocols, 1997. Proceedings., 1997 International Conference on , 28-31 Oct. 1997

Page(s): 95 -104

[\[Abstract\]](#) [\[PDF Full-Text \(920 KB\)\]](#) **IEEE CNF****2 LCM: a multicast core management protocol for link-state routing networks***Yih Huang; Fleury, E.; McKinley, P.K.;*

Communications, 1998. ICC 98. Conference Record.1998 IEEE International Conference on , Volume: 2 , 7-11 June 1998

Page(s): 1197 -1201 vol.2

[\[Abstract\]](#) [\[PDF Full-Text \(512 KB\)\]](#) **IEEE CNF****3 Reliable multicast in multi-access wireless LANs***Kuri, J.; Kasera, S.K.;*

INFOCOM '99. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE , Volume: 2 , 21-25 March 1999

Page(s): 760 -767 vol.2

[\[Abstract\]](#) [\[PDF Full-Text \(816 KB\)\]](#) **IEEE CNF****4 Mobility based dynamic multicast protocol in wireless ATM network**

Gyung-Ho Hwang; Dong-Ho Cho;

Vehicular Technology Conference, 1999. VTC 1999 - Fall. IEEE VTS

50th , Volume: 4 , 19-22 Sept. 1999

Page(s): 2009 -2013 vol.4

[\[Abstract\]](#) [\[PDF Full-Text \(432 KB\)\]](#) **IEEE CNF**

5 The wave expansion approach to broadcasting in multihop radio networks

Chlamtac, I.;

Communications, IEEE Transactions on , Volume: 39 Issue: 3 , March 1991

Page(s): 426 -433

[\[Abstract\]](#) [\[PDF Full-Text \(708 KB\)\]](#) **IEEE JNL**

[Home](#) | [Log-out](#) | [Journals](#) | [Conference Proceedings](#) | [Standards](#) | [Search by Author](#) | [Basic Search](#) | [Advanced Search](#)
[Join IEEE](#) | [Web Account](#) | [New this week](#) | [OPAC Linking Information](#) | [Your Feedback](#) | [Technical Support](#) | [Email Alerting](#)
[No Robots Please](#) | [Release Notes](#) | [IEEE Online Publications](#) | [Help](#) | [FAQ](#) | [Terms](#) | [Back to Top](#)

Copyright © 2003 IEEE — All rights reserved

A Highly Available Local Leader Election Service

Christof Fetzer, Flaviu Cristian

Abstract—We define the highly available local leader election problem, a generalization of the leader election problem for partitionable systems. We propose a protocol that solves the problem efficiently and give some performance measurements of our implementation. The local leader election service has been proven useful in the design and implementation of several fail-aware services for partitionable systems.

Keywords—Local leader election, partitionable systems, timed asynchronous systems, global leader election.

I. INTRODUCTION

THE leader election problem [1] requires that a unique leader be elected from a given set of processes. The problem has been widely studied in the research community [2], [3], [4], [5], [6]. One reason for this wide interest is that many distributed protocols need an election protocol as a sub-protocol. For example, in an atomic broadcast protocol the processes could elect a leader that orders the broadcasts so that all correct processes deliver broadcast messages in the same order. The *highly available leader election problem* was defined in [7] as follows: (S) at any point in time there exists at most one leader, and (T) when there is no leader at time s , then within at most κ time units a new leader is elected.

The highly available leader election service was first defined for synchronous systems in which all correct processes are *connected*, that is, can communicate with each other in a timely manner. Recently, the research in fault-tolerant systems has been investigating asynchronous partitionable systems [8], [9], i.e. distributed systems in which the set of processes can split in disjoint subsets due to network failures or excessive performance failures (i.e. processes or messages are not timely; see Section III for details). Like many other authors do, we call each such subset a *partition*. For example, processes that run in different LANs can become partitioned when the bridge or the network that connects the LANs fails or is “too slow” (see Figure 4). One reason for the research in partitionable systems is that the “primary partition” approaches [10] allow only the processes in one partition to make progress. To increase the availability of services, one often wants services to make progress in all partitions.

Our recent design of a membership [11] and a clock synchronization service for partitionable systems [12] has indicated that we need a leader election service with different

properties for partitionable systems than for synchronous systems. The first problem that we encountered is how to specify the requirements of such a *local leader election service*. Ideally, such a service should elect exactly one local leader in each partition. However, it is not always possible to elect a leader in each partition. For example, when the processes in a partition suffer excessive performance failures, one cannot enforce that there exists exactly one local leader in that partition. To approach this problem we have to define in what partitions local leaders have to be elected: we introduce therefore the notion of a *stable partition*. Informally, all processes in a stable partition are connected to each other, i.e. any two processes in a stable partition can communicate with each other in a timely manner. The processes in a stable partition are required to elect a local leader within a bounded amount of time. An election service might be able to elect a local leader in an *unstable partition*, i.e. a partition that is not stable, but it is not guaranteed that there will be a local leader in an unstable partition. We call a process “unstable” when it is part of an unstable partition.

In each stable partition, a local leader election service has to elect exactly one local leader. In an unstable partition the service might not be able to elect exactly one local leader. It can be advantageous to split an unstable partition into two or more “logical partitions” with one local leader each if that enables the processes in each of these logical partitions to communicate with each other in a timely manner (see Figure 1). To explain this, note that our definition of a “stable partition” will require that processes in such a partition be connected to each other. This implies that when the connected relation in a partition is not transitive, that partition is unstable. For example, the connected relation can become non-transitive for three processes $\{p, q, r\}$ if the network link between p and r fails or is overloaded while the links between p and q and q and r stay correct (see Figure 2).

In specific circumstances, our local leader service splits an unstable partition into two or more logical partitions with one leader in each. The service makes sure that timely communication between any two processes in a logical partition is possible. However, sometimes this communication has to go via the local leader in case two processes p and r in a logical partition are only connected through the local leader q (see Figure 2.b). Informally, a logical partition is a set of processes that can communicate with each other in a timely manner through a single local leader.